AN IMAGING DEVICE THAT USES THE WAVELET TRANSFORMATION AS THE IMAGE RECONSTRUCTION ALGORITHM. III. DIGITAL FILTERING OF IMPERFECT WAVELETS.

J. H. Letcher, K. A. O'Neil, and P. J. Cook

University of Tulsa, Tulsa, Oklahoma USA

Abstract. An ultrasound imaging device that uses the wavelet transformation as the image reconstruction algorithm has been described in the first papers of this series. In this paper we describe a digital filtering technique which improves the practicality of the device by easing the constraints on sound generation, allowing the use of simple and inexpensive sound transducers and drivers. The extent to which filtering can compensate for poor control of the generated sound is investigated through simulation and experiment.

INTRODUCTION

The previous papers in this series describe an ultrasound imaging device with an image reconstruction algorithm based on the wavelet transform [Letcher 3,4]. Briefly, the sound pulse directed into the sample has the shape of a "wavelet." The echo returned from the sample is transformed to reconstruct the acoustic impedance of the sample along the path of the sound pulse. The effectiveness of the algorithm depends on the accurate production of sound pulses having the correct shape, one which the transducer is not naturally disposed to produce. The approach taken in the work cited was to analyze the impulse response of the transducer and pre-distort the transducer's drive signal to get the correct output. The implementation requires high-speed, high-voltage amplifiers, which are costly.

In this paper we investigate a different approach that requires only the simplest hardware -- an inexpensive transducer and an impulse driver. The sound pulse produced is a distorted wavelet, and the echo obtained from the sample is likewise distorted. The echo signal is digitally processed to remove the distortion and produce the echo signal that would have been received had the transmitted sound pulse been a faithful reproduction of a wavelet. As a result, we can use the wavelet transform to reconstruct the image while reducing the cost and complexity of the hardware.

The organization of the paper is as follows: In the section below we describe the mathematical model of the experiment and the requirements for the digital filter which removes the distortion. The next sections contain filter design algorithms and criteria for the stability and performance of the filters. A recursive filter is used in a computer simulation, and non-recursive filters are designed to correct for distortion in an actual experiment.

MATHEMATICAL MODEL OF THE EXPERIMENT

We will concentrate on a one-dimensional ultrasound experiment which measures the acoustic impedance of the sample along a line. A focused pulse of sound is directed into a sample and the reflected sound signal is detected by a transducer and recorded. Sound is reflected from the interface between materials of differing acoustic impedance. The echo signal is a superposition of delayed versions of the transmitted pulse; the amplitude, phase, and delay of each reflection are related to the nature and position of the corresponding interface.

In our mathematical model of the experiment we shall assume that all reflections are due to weak discontinuities in impedance (on the order of a few percent), and hence that there are no multiple reflections. We also assume that the sound velocity in the sample is constant and that there is constant absorption in the sample. These assumptions are reasonable for a biological sample. We also assume that the interfaces are discrete, and hence we model continuous (extended) interfaces by a sequence of weak discrete ones.

All signals are treated digitally, so we express signals as sequences of samples of the analog signal taken at regular intervals. We will assume that the sampling rate is high enough to clearly resolve the sound signal. Sequences are denoted by bold face letters, and the elements by subscripted letters. We use the convention that all elements with negative subscripts are zero and that finite length sequences may be padded by zeros to infinite length in both directions.

The wavelet transform is an algorithm for approximating a function by a superposition of functions having a "model" shape, called the wavelet. These basis functions are all related by affine transformations of the time variable, which is to say that their graphs differ only by horizontal translations and scalings. If the wavelet shape is carefully chosen, the basis functions are mutually orthogonal (the integral of the product vanishes) and the transform is "fast," i.e. the amount of computation required to transform a sequence of length N is proportional to $N \log N$. The result of the transformation is the sequence of coefficients of the wavelet basis functions in the approximation. We will use a non-orthogonal basis of wavelets, described elsewhere [Letcher 4]. The transform with respect to this basis is related to the transform with respect to the orthogonal basis by multiplication by a sparse matrix, and hence is also "fast."

Let $s = (s_0, \ldots, s_m)$ be the transmitted signal and $r = (\ldots,0,r_0,r_1,r_2,\ldots)$ be the received sound signal. The acoustic impedance of the sample along the line probed by the sound pulse is characterized by a sequence **x**, which relates the transmitted sound to the echo by [Letcher 3]:

$$r_i = \sum_j s_j x_{i-j} \quad \text{(1a)}$$

that is,

$$r = s * x \quad \text{(1b)}$$

where "*" is the convolution operator, defined by Equation (1a).

The wavelet transform using the non-orthogonal basis generated by the wavelet s is in effect a deconvolution $(s*)^{-1}$ which recovers **x** from **r**. The acoustic impedance of the sample can be computed from the sequence **x**. If $W(\cdot)$ and $W_L(\cdot)$ denote the wavelet transform with respect to the orthogonal and non-orthogonal bases, respectively, then we have $\mathbf{x} = W_L(\mathbf{s}*\mathbf{x}) = W_L(\mathbf{r})$.

In general, one expects that the transmitted sound sequence s will not be a wavelet. We shall think of s as a distorted version of a wavelet sequence **w**. The advantage of this point of view is that we never have to generate a sound pulse shaped like a wavelet. We will assume that there is a (reasonably short) distorting sequence **d** so that

$$s = d * w . \quad \text{(2)}$$

This relation works well for the inexpensive transducers that we used in our experiment. Henceforth all wavelet transforms will be taken with respect to the bases generated by **w**.

In terms of **w**, the received signal is $\mathbf{r} = \mathbf{d}*\mathbf{w}*\mathbf{x}$. We propose to take the distortion into account by convolving r with a filter sequence **f** that has the effect of the deconvolution $(\mathbf{d}*)^{-1}$. Subsequent wavelet transformation will then recover the sequence **x** with no distortion.

FILTER DESIGN

Given the sequence **s**, we must find sequences **w** and **f** that either exactly or approximately satisfy the relation $\mathbf{f} * \mathbf{s} = \mathbf{w}$ and so that **w** generates a wavelet function. Equivalently, we also find a sequence **d** so that the relations $s = \mathbf{d} * \mathbf{w}$ and $\mathbf{f} * \mathbf{d} = \mathbf{I}$ hold, where **I** is the convolution identity $(\ldots,0,0,1,0,0,\ldots)$.

The basic mathematical problem is to recover the sequence **u** given data $\mathbf{d} = (d_0, \ldots, d_m)$ and $\mathbf{v} = \mathbf{d} * \mathbf{u}$. Solving for u_i in the relation

$$v_i = \sum_{j=0}^{m} d_j u_{i-j} \quad \text{(3)}$$

yields the recursive relations

$$u_i = (1/d_0)\,(v_i - \sum_{j=1}^{m} d_j u_{i-j}). \quad \text{(4)}$$

These are recursive because u_i cannot be computed without first computing u_{i-1}. Equation (4) describes the action of a recursive filter, and is equivalent to convolving **v** with an infinite-length sequence. We shall denote this operation by $(\mathbf{d}*)^{-1}$.

In general there is no sequence **f** of finite length so that $\mathbf{f}*\mathbf{v} = \mathbf{u}$, or equivalently $\mathbf{f}*\mathbf{d} = \mathbf{d}*\mathbf{f} = \mathbf{I}$, is satisfied exactly. A sequence that approximately satisfies these relations will be denoted $\mathbf{d}^{-1}$. To support our convention that elements with negative indices are zero, we will when necessary substitute for **I** a sequence with only one nonzero entry, occurring at some non-negative index l, called the lag. The sequence $\mathbf{d}^{-1}$ is also known as a spiking filter [Robinson and Trietel 1]. While this "inverse" notation is convenient, one must remember that $\mathbf{d}^{-1}$ is not unique. However, given a fixed length for the inverse sequence and a fixed lag, there is a unique sequence which best inverts convolution with **d** in the least-squares sense. A simple formula for these spiking filters is found in [Robinson and Trieitel 1]:

$$d^{-1} = A^{-1} \cdot g, \quad \text{(5)}$$

where

$$A_{ij} = \sum_k d_k d_{k-i+j} \quad \text{(6)}$$

and

$$g_i = d_{l-i} \quad \text{(7)}$$

Given these formulas for $(\mathbf{d}*)^{-1}$ and $\mathbf{d}^{-1}$, there are several algorithms for recovering **x** from **r** and **s**. The first step is to pick a wavelet **w** that generates an orthogonal basis for the wavelet transform W, so that W_L is fast and $\mathbf{x} = W_L(\mathbf{w} * \mathbf{x})$. To obtain the sequence $\mathbf{w} * \mathbf{x}$, one may first calculate **d** as $W_L(\mathbf{s})$, $(\mathbf{w}*)^{-1}(\mathbf{s})$, or $\mathbf{w}^{-1} * \mathbf{s}$; then compute either $(\mathbf{d}*)^{-1}(\mathbf{r})$ or $\mathbf{d}^{-1} * \mathbf{r}$. The intermediate result **d** may be avoided by computing the deconvolution of **d** as either $\mathbf{w} * (\mathbf{s}*)^{-1}$ or $\mathbf{w} * \mathbf{s}^{-1}$. Finally, one may use linear combinations of spiking filters to construct a filter **f** that approximately satisfies $\mathbf{f} * \mathbf{s} = \mathbf{w}$, which is in effect $\mathbf{d}^{-1}$. This approach tends to be the most straightforward one.

EFFECTS OF FILTERING

While the recursive filter is an exact deconvolution, it cannot always be used. Implementing the recursive filter $(\mathbf{d}*)^{-1}$ is equivalent to solving a nonhomogeneous difference equation

$$v = (\sum_{j=0}^{m} d_j B^j)\, u \qquad (8)$$

where B is the "backshift operator" $(B\mathbf{u})_i = u_{i-1}$. The homogeneous equation $0 = (\Sigma d_j B^j)\mathbf{u}$ has exponentially growing solutions if the characteristic polynomial $P(t) = \Sigma d_j\, t^j$ has any nonzero root inside the unit circle [Fuller 2]. These solutions give rise to numerical instabilities which render the recursive filter impractical in these cases. Thus the recursive filter cannot be used unless **w** can be chosen so as to give P(t) the desired properties.

The non-recursive convolution filter is stable, but gives only an approximate deconvolution. The general effect of a well-designed least-squares filter is that of a deconvolution with "ripples" added. This is not stochastic noise, but rather a reflection of the fact that $(\mathbf{d}^{-1} * \mathbf{d}) - \mathbf{I}$ is not identically zero. For the signals encountered in our ultrasound experiments, with filter length equal to the wavelet length, the ripple amplitude was roughly ten percent of the signal amplitude. This amplitude decreases with increasing filter length, and also depends on the wavelet shape.

Although the mathematical model assumes that the deconvolution filter will be applied only to sequences that result from convolution, in practice there will be noise in the signal that does not have this origin. It is therefore of value to know the effect of the filter on noise. This is described by the squared gain of the filter.

Suppose **u** is a mean zero covariance-stationary random sequence, E is expectation, and $\mathbf{v} = \mathbf{u} * \mathbf{d}$. We define the autocovariance sequence of **u**, γ_u, to be

$$(\gamma_u)_j = E(\sum_k u_k u_{k+j}) \qquad (9)$$

and the spectrum of **u**, $S_u(f)$ to be

$$S_u(f) = \sum_k (\gamma_u)_k \exp[-2\pi i k f],\ -\frac{1}{2} \le f \le \frac{1}{2}. \qquad (10)$$

f is the frequency in units of 1/T where T is the interval between samples. The spectrum $S_u(f)$ can be interpreted as a decomposition of the variance of **u**, because the spectrum is large over those frequency bands which account for a large fraction of the variation of **u**. Then the ratio $S_v(f)/S_u(f)$ indicates the amplification of the convolution filter **d** at the frequency f. This ratio, called the squared gain of the filter **d** at frequency f, can be computed directly from the sequence **d**:

$$|A_d(f)|^2 = |\sum_k d_k \exp[-2\pi i k f]|^2. \qquad (11)$$

COMPUTER SIMULATION OF EXPERIMENT

To verify the algorithm, we performed a computer simulation of the ultrasound experiment. The results are depicted in Figure 1. The sequence **x** was chosen to correspond to the sample impedance given in Fig. 1a and the transmitted sound is shown in Fig. 1b. The Haar wavelet was selected as the basis for the wavelet transform. Without filtering, the reconstructed image of sample impedance shows considerable degradation (Fig. 1). We implemented the recursive filter $(\mathbf{d}^*)^{-1}$, which was stable, before the wavelet transform to produce the image in Fig. 1d. The distortion is clearly reduced by the filtering. A noisy experiment was simulated by adding a noise signal to the received signal. The noise amplitude was one-tenth the signal amplitude. The reconstructed impedance is shown in Fig 1e.

DESCRIPTION OF EXPERIMENTAL APPARATUS

An apparatus was assembled to verify the performance of the reconstruction algorithm on experimental data. A "phantom" constructed from parallel blocks of wax 6.35 mm thick held approximately 25 mm apart by spacers was immersed in a beaker of water. A Panametrics 2.5 MHz V305SU F=3" ultrasonic transducer, connected to a Panametrics 500PR pulser-receiver, was aimed at right angles to the wax layers so that echoes produced at the wax-water interface would return to the transducer.

The received ultrasound signal was led to a custom circuit board which, when triggered by a sync pulse from the pulser, amplified the signal, sampled it at a 20MHz rate, performed analog-to-digital conversion, and stored the resulting data in a cache memory. The analog portion of the circuitry is identical to that on the Sony CXA 1296P PCB evaluation board, built around a Sony CXA 1296P eight-bit analog-to-digital converter. The data was later uploaded to a PC-compatible computer equipped with a Cyrix 486 DLC processor running at 33MHz, a Cyrix EMC 87 coprocessor, an ATI VGA Wonder 512 graphics adaptor and a NEC 5FG monitor.

As a result of less-than-complete shielding, noisy power supplies, excessive cable lengths, etc., the experimental data contained noise. The noise was reduced by averaging the data from 350 separate pulses.

EXPERIMENTAL RESULTS

The results of the digital filtering algorithm on the reduced-noise data set are shown in Fig. 2. The gain of the analog stage before the A/D converter was adjusted to make the first echo amplitude nearly full-scale. Because the interfaces in the phantom were sharp, each echo received should have had the same shape as the transmitted sound signal. Figure 2a shows a superposition of the six echoes received, rescaled to have the same amplitude. The interval between samples is 50 nsec. The variations between the echoes can be taken as a measure of the total noise in the experiment. One of the echoes was used as **s**, the transmitted sound signal, in the design of a non-recursive convolution filter of length 60. The echo after filtering, shown in Fig. 2b, is seen to be a Haar wavelet with very little distortion. Figure 2c shows the complete echo sequence after filtering and wavelet transformation. The decreasing amplitude of the echoes is due to the absorption in the wax, and could easily be accounted for in the image reconstruction algorithm. The reconstructed acoustic impedance of the sample (without correcting for absorption) is shown in Fig. 2d. Note that the noise in the signal has an effect on the performance of the algorithm. The squared gain of the filter used, shown in Fig. 2e, indicated considerable amplification at high frequencies.

The same filter was then applied to the data without the benefit of averaging to reduce the noise. The result, seen in Figure 3, is significantly degraded and clearly shows the amplified high frequency components of the noise.

SUMMARY AND CONCLUSIONS

In this paper we have described a modification of the reconstruction algorithm based on the wavelet transform which does not require that the generated sound have the shape of a wavelet. The essential feature of the modification is the implementation of a deconvolution filter which in effect converts the received sound signal into a superposition of wavelets. The advantage of this approach is that the numerically efficient wavelet transform is retained as the basis for the reconstruction algorithm, but the requirement for high-accuracy sound generation is eased.

The deconvolution filter may be implemented recursively or non-recursively. The recursive filter gives excellent results and is very efficient numerically, but suffers from instability if the corresponding characteristic polynomial has unstable roots (lying outside the unit circle). Generally speaking, if the peaks of the sound signal are decreasing in amplitude, the recursive filter will be stable.

The non-recursive filter is easy to compute, reasonably fast, always stable, and gives excellent results on noiseless signals. However, the filters tend to be sensitive to noise, and to amplify the higher frequencies of the noise in particular. In theory, a well designed medical imager can have a very low noise level, since there is no naturally occurring noise in biological tissue at these frequencies. The algorithm described in this paper would work very well for such a device.

References:

1. Robinson,E. and Trietel,S.,1980, Geophysical Signal Analysis, Prentice Hall.
2. Fuller,W.A.,1976, Intro. to Statistical Time Series, John Wiley,46.
3. Letcher,J.H.,1992, Int.J. Imag. Sys. and Tech., 27, 98-108.
4. Letcher,J.H.,1993, Int.J. Imag. Sys. and Tech., submitted for publication.

Figure 1a: Simulated sample impedance

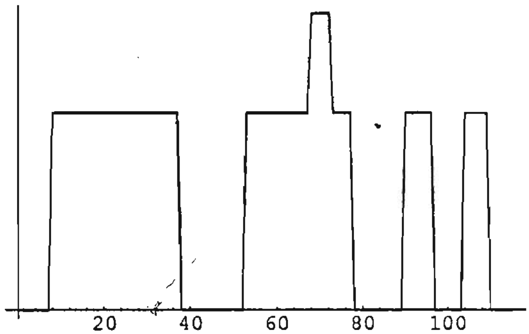

Figure 1b: Simulated transmitted sound

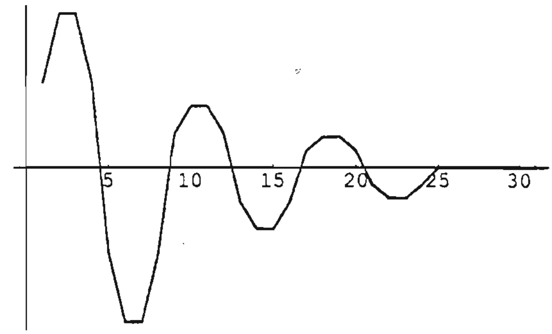

Fig. 1c: Reconstruction (no filtering)

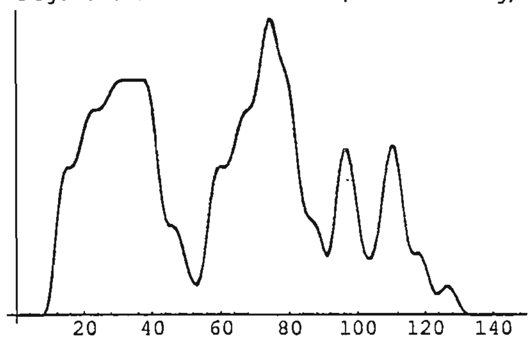

Fig.1d: Reconstruction with filtering

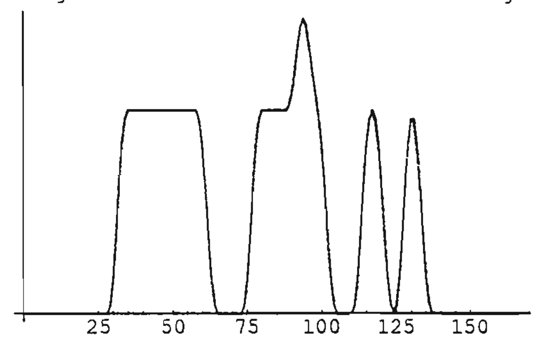

Fig.1e: Reconstruction of noisy signal

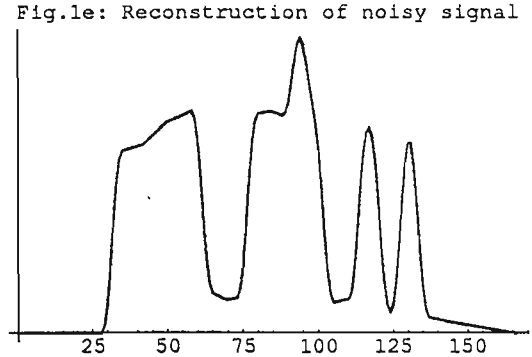

Figure 2a: Superposition of echoes

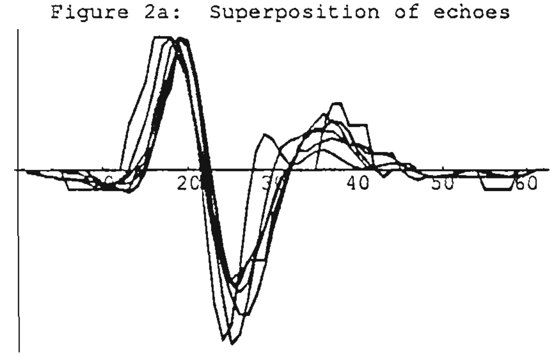

Figure 2b: Echo after filtering

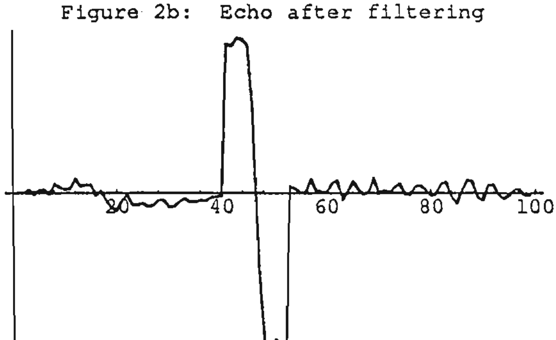

Figure 2e: Squared gain of filter

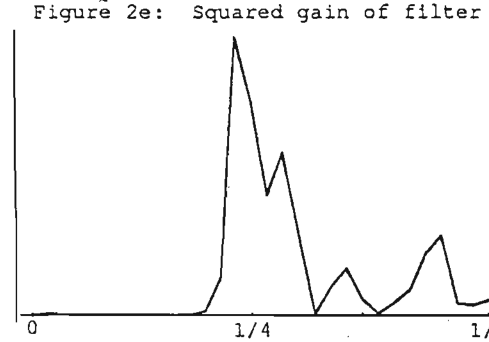

Figure 2c: Reduced-noise data after filtering and wavelet transform

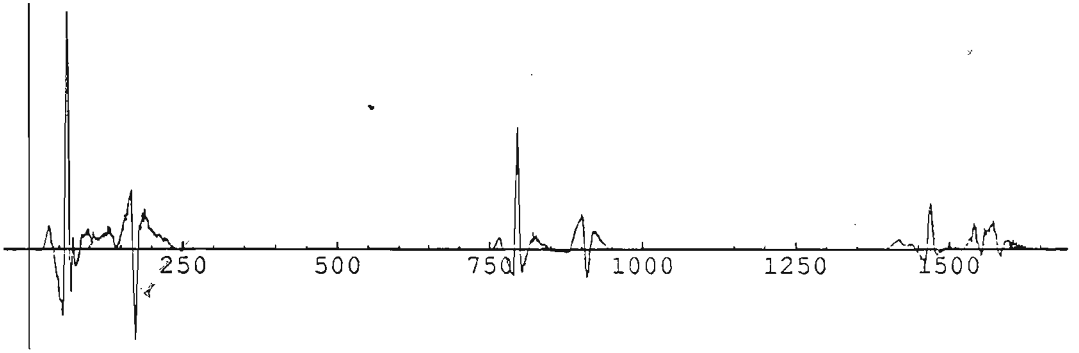

Figure 2d: Reconstructed sample impedance

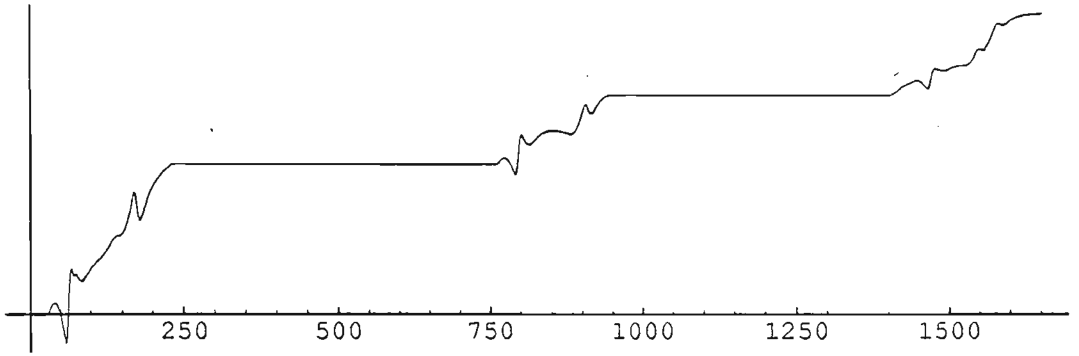

Figure 3: Filtered noisy data